Proc IEEE Computer Society 18th Conference on Local Computer Networks, Minneapolis, MN, September 19-22, 1993. This pdf reconstructed Dec 2005.

Using the Common LAN to Introduce ATM Connectivity.

G. J. Armitage (gja@ee.mu.OZ.AU)
K.M. Adams (kma@ee.mu.OZ.AU)

Deptartment of Electrical and Electronic Engineering,
University of Melbourne, Parkville, 3052, Australia.

Abstract.

This paper outlines a method for using LAN technologies to transport Asynchronous Transfer Mode (ATM) cells. B-ISDN service requirements are broken into three groups - high speed media, multi-media interfaces, and service control software. Compression techniques for bandwidth intensive services are overviewed. It is argued that these services may be provided using sub-Mbit/sec ATM links. The ATM concept is shown to be independent of the physical layer, allowing low-speed services to be implemented and utilised before high speed links are in place. An architecture is described where a LAN/fibre gateway switches cells between an external B-ISDN fibre and an in-house Ethernet based ATM network.

1. Introduction.

One of the significant driving forces of Broadband ISDN is the desire to implement universal multi-media communication - integrating voice, video, still pictures, and digital data. The ITU Telecommunication Standardisation Sector (ITU-TSS, formerly the CCITT) has adopted Asynchronous Transfer Mode (ATM) to provide flexible and uniform physical layer transport [8, 9]. Many research goups are working on the integration of various media at the user interface, developing software and hardware to combine ATM networks with advanced multi-media user interfaces.

An issue which is now gaining prominence is how this technology can be introduced into the network market? 600Mbit/sec fibre technology is becoming yesterdays breakthrough, and Gbit speeds are being promised. However the average business or personal users are only just beginning to wonder how they can introduce this leading edge technology. The answer we propose is to apply some of today's bread and butter technology as the stepping stones to full ATM implementations.

The upgrade of an office or business to full multimedia B-ISDN connectivity requires three broad areas of change. Optical fibre cables will be laid to provide the high bitrates that B-ISDN is intended to support. New or upgraded workstations containing the appropriate codecs will be required to support real-time video and audio connections. Finally new software will be required to allow users to connect, for example, to B-ISDN based image databases or interactive multi-media document servers.

The key decision we propose is to begin by creating a hybrid ATM network of fibre and copper based LAN technologies within the premises. ATM traffic is carried in LAN packets between the desktop and a gateway where the B-ISDN fibre enters the customer's premises. Hardware upgrades are minimal, with the focus being on providing present generation workstations and high end PC's with software based ATM Adaptation Layers (AALs). We also look at the way advances in image compression methods are reducing the bit rate needs of multi-media applications. The practical limitations of our scheme will be discussed.

Section 2 looks at the basic components of multmedia, and the advances in image compression technology. ATM is introduced in section 3 and described briefly. In section 4 the basic components of a LAN based ATM network are described. Some implementation possibilities are described in section 5, and section 6 touches on the sorts of performance limits we can expect. The paper is concluded in section 7.

2. Basic multi-media services.

Today's desktop computer is capable of displaying text, graphical information, and real time moving images. In addition, many already have digitised sound capabilities built in. Unfortunately real-time services such as sound and video have needed orders of magnitude more network bandwidth than traditional services such as file transfers and inter-node signalling. This has made video/voice conferences appear impractical over conventional LAN technologies such as Ethernet, Token Ring, and LocalTalk. 'Conventional wisdom' has also claimed that LANs couldn't offer sufficiently predictable performance to satisfy such delay-sensitive services. The transfer of static color images across the network also imposes a significant peak load, as raw bit mapped images frequently consume hundreds and thousands of kbytes. Some studies have concluded that users will expect systems to respond with images within 2 seconds of a selection being made

[1]. Rapid retrieval and display of such images places a greater peak load on the network than more traditional text retrieval.

2.1 Still images

A 24 bit 'true color' NTSC frame sized picture (480 by 640 pixels) consumes 900 Kbytes. Bringing a more reasonable 960 by 1280 pixel 'true colour' picture to the screen within 2 seconds would require network traffic in excess of 14.4 Mbit/sec (ignoring processing time).

Various schemes are being developed for encoding and storing still images. The Graphical Interchange Format (GIF) offers exact reproduction of an image using non-lossy encoding techniques. For example, an 1152 by 800 pixel, 8 bit per pixel colour image consumed only 470 KB in GIF format (the exact compression will depend on the source material).

The ISO/CCITT Joint Photographic Experts Group is developing a lossy image storage algorithm called JPEG, intended for the storage of photo quality "real world" images. The JPEG encoding algorithm allows a user to specify the acceptable "lossiness" on a per image basis, with decoding being independent of the encoded loss level. Our 470 KB GIF file consumed only 90KB when JPEG encoded at a quality value of 75% (using the 'xv' image viewer written by John Bradly at University of Pennsylvania). At a quality value of 50% the image consumed only 57KB. Subjective assesment of the images at 1 metre from the screen suggested very little degradation could be seen at the 75% level. At the 50% level the picture was still quite acceptable, although the image brightness was starting to look exaggerated.

In [2] the JPEG standard is described as producing almost perfect images with a compression ratio of 5:1, and moderate picture quality with compression as high as 30:1. JPEG encoding and decoding hardware is becoming available for current workstations. This suggests that retrieving and displaying JPEG encoded still images within 2 seconds will be possible, while using significantly less network bandwidth than raw bitmaps.

2.2 Moving images

The ways in which video and audio media may be used are varied. Some people conceive of video conferences where approximately life size images are distributed across the digital network. The VideoWindow system developed by BellCore is one such scheme, using 3ft by 8ft display screens at each end of a link. Research with VideoWindow revealed that public video conferencing was relatively insensitive to variations in transmission bandwidth ranging from 384kbit/sec to 45 Mbit/sec [3].

Another image is of small video windows sharing real estate with other windows on the traditional workstation or PC screen. A decrease in acceptable image size decreases the required image resolution. This has implications for the types of video compression and encoding algorithms that may be profitably used. The ISO Moving Pictures Experts Group (MPEG) has developed an encoding scheme to allow the delivery of real time VCR quality video and audio within the constraints of 1.544Mbit/sec data links [4]. Currently it demands intensive processing at the encoding end, but decoding is easily achieved in real time. MPEG appears well suited for Video On Demand video distribution systems [5].

Another encoding scheme of interest is CCITT's H.261. This provides digital video at rates between 64 kbit/sec to 1920 kbit/sec (in increments of 64 kbit/sec). At 64 and 128 kbit/sec H.261 is considered acceptable for videophone style applications (at 176 by 144 pixels). 384 kbit/sec is considered as a minimum to reasonably support a video conference (at 352 by 288 pixels)[6]. The image quality increases with bit rate, but "has been perceived as less than VCR quality at 1.544 Mbit/sec" [5]. H.261 is also notable in that its encoding involves substantially less processing than MPEG. Video is already appearing on the Internet, with software codecs (using H.261 based algorithms) transmitting video within User Datagram Protocol (UDP) packets [7].

2.3 Practical consequences.

Developments such as JPEG, MPEG, and H.261 are reducing the network capacity needed to support fairly basic image and videophone applications. Careful structuring of interactive multimedia documents can also reduce the peak bitrates needed by spreading out the intervals between user requests for new information [24].

The speed penalty inherent in ATM-over-LAN appears tolerable for low quality video and image applications. It makes sense to develop ATM based multi-media applications and network interfaces even before fibre optic physical layers are in place, and with suitable physical layer gateways gain access to ATM based services in the rest of the B-ISDN.

3. Asynchronous Transfer Mode.

ATM is a packet switched data transport system based on fixed size 53 byte cells - 48 bytes of user data and a 5 byte header (Figure 1). Each cell carries a Virtual Channel Indicator (VCI) and Virtual Path Indicator (VPI) in its header. Cells are routed through switching nodes, using the combination of VPI and VCI as a label to associate each cell with established virtual connections.

Proc IEEE Computer Society 18th Conference on Local Computer Networks, Minneapolis, MN, September 19-22, 1993. This pdf reconstructed Dec 2005.

Essential to the services offered by the new ATM networks is the ATM Adaptation Layer (AAL). The AAL is a CCITT defined layer which "adapts" the cell based ATM physical layer to packet, datagram, or bitstream oriented higher layers [10, 11]. It exists at the endpoints of virtual connections, where higher layer peers wish to establish communication.

3.1 VPIs, VCIs, virtual connections.

In section 2 of CCITT Recommendation I.150 [8] ATM is described as "*a connection-oriented technique. Connection identifiers are assigned to each link of a connection when required, and released when no longer needed.*". The VPI and VCI fields within each cell header are the connection identifiers. In section 3.1.1 of I.150 an ATM connection is described as consisting "*..of the concatenation of ATM layer links in order to provide an end-to-end transfer capability to endpoints.*". An ATM link consists of any cell path between points where the VPI or VCI are switched or terminated.

The VCI field is 16 bits wide. The VPI field has two CCITT defined sizes. At the Network Node Interface (NNI) it is 12 bits, and at the User Network Interface (UNI) it is 8 (the remaining 4 bits carrying Generic Flow Control (GFC) information) [17]. CCITT have established an hierarchical relationship between VPIs, VCIs, and the physical media [18]. Virtual channels are considered to exist within virtual paths, which are themselves unique only within a given physical path. Figure 2 illustrates the basic relationships.

Virtual Channel Connections (VCCs) provide end-to-end cell paths between users of the ATM layer (typically the AALs at either end). We shall simply refer to them as Virtual Connections for the rest of this paper. Both the VPI and VCI values may change along a VCC on a per-link basis. Virtual Path Connections (VPCs) carry bundles of VCCs across sequences of links where only the VPI is changed. Within a VPC, and hence a VCC, cell

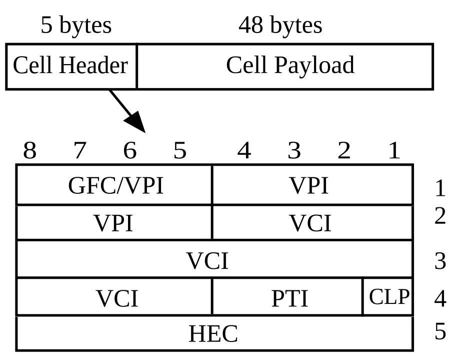

Figure 1

sequence is assured. The physical media are responsible for carrying ATM cells between physical nodes. A node may provide VPI or VPI/VCI based cell switching, or it may terminate the virtual connection at an AAL.

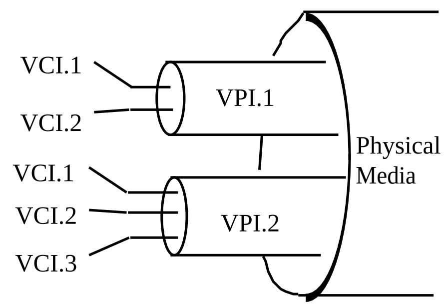

Figure 2

Figure 3 represents a virtual connection spanning three switching nodes and four physical links. At node 1 the VCI is switched, at node 2 both the VPI and VCI are switched, and at node 3 the cells are switched to another physical link without changing the VPI or VCI. The allocation of VPI and VCI values to each link of a virtual connection is the responsibility of the signalling protocol in use.

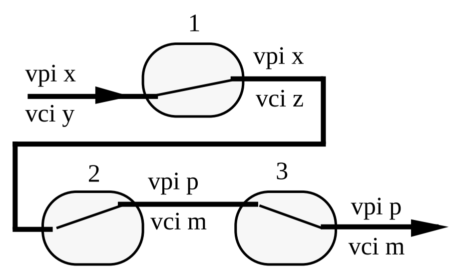

Figure 3

3.2 The ATM Adaptation layer (AAL).

A virtual connection end point consists of a physical medium interface, an ATM layer, and an AAL (Figure 4). The ATM layer demultiplexes received cells based on their VPI and VCI before passing them upwards to the AAL. Cells which do not belong to any virtual connection currently open to the AAL are filtered out and dropped by the ATM layer. An overview of the AAL's functions will be given here. For a more detailed description CCITT recommendations I.362 and I.363 [10, 11] should be referred to.

Figure 4 shows multiple user services being simultaneously supported by the AAL. Each service may require different data transfer mechanisms - CCITT is defining five different AALs to provide the necessary range of functions. An example is having simultaneous video and file transfers across a single physical medium. Each service would have its own virtual connection, with different AALs above the ATM layer being the termination point of each virtual connection.

The AAL is considered to consist of two main sublayers: the Convergence Sublayer (CS) and the Segmentation and Reassembly sublayer (SAR) . The CS function provides any necessary encapsulation of user data before the SAR function inserts the result into 48 byte cell payload fields. The SAR function provides transmission and error detection facilities on a cell by cell basis. The CS function provides transmission and error detection facilities over the 'natural' unit of data utilised by the user service (bytes, bitstreams, or variable length packets). In some applications the SAR or CS functions may be empty.

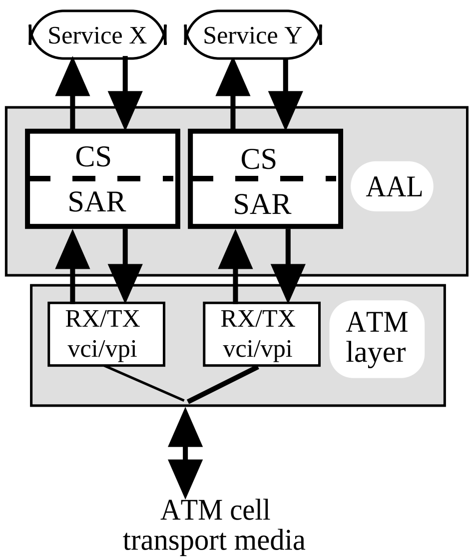

Figure 4

For constant bitrate services, section 2 of I.363 [11] specifies AAL type 1 (AAL1). For continuously variable bitrate services, AAL type 2 is being defined (AAL2, section 3 of I.363). AAL1 is reasonably well defined at this stage, and is expected to be used by digitised video and voice services. These services require clock signals to be provided by the AAL, and some measure of buffering within the AAL, to ensure that information crosses the AAL-higher layer boundary at a fixed rate.

Data communication protocols (such as the DARPA Internetworking Protocol , IP [15, 16]) are based around the arbitrary and unpredictable transmission of packets or frames of information. Support for packet oriented services is presently provided by two 'standard' AALs, Type 3/4 (AAL3/4, section 4 of I.363), and Type 5 (AAL5, [12]). AAL3/4 is the result of merging two earlier AALs, type 3 and type 4. AAL5 is a product of the datacomms industry desire for a packet oriented AAL which maximised the efficiency of cell utilisation by using as little encapsulation as possible [13, 14]. The evolving draft standard for IP over ATM proposes to use AAL5 [26].

The AAL hides the various B-ISDN services from the ATM layer. The ATM layer hides the AAL from the physical cell transport layer. The only characteristic required of a physical layer is that it transport cells between the ATM layers of user endpoints and the ports of cell switching nodes. The rest of this paper looks at the use of LAN technology to transport ATM cells.

4. A unique in-house ATM network.

A practical ATM network is going to consist of physical links between switching nodes and AAL endpoints (service users such as workstations, servers, and multi protocol network layer gateways or routers). For most business or private users a single physical link will be provided between their premises and the local B-ISDN "exchange". An in-house switching node will behave like an ATM PABX, allowing virtual connections to be established between in-house ATM based services and remote services.

The ATM concept outlined in section 3 is independent of the nature of the physical medium used to carry cells between nodes. The use of optical fibre has been pushed because the B-ISDN vision encompasses bandwidth intensive services. However as identified in section 2, many basic multi-media services may be implemented with quite low bandwidth requirements. Figure 5 shows a new arrangement where the in-house ATM network is a hybrid of "ATM fibre" and Mbit/sec LAN (e.g Ethernet [19]). The key components of this new network are the fibre/LAN cell switching node, and workstations equipped with software ATM layers and AALs on top of their pre-existent LAN interfaces.

Proc IEEE Computer Society 18th Conference on Local Computer Networks, Minneapolis, MN, September 19-22, 1993. This pdf reconstructed Dec 2005.

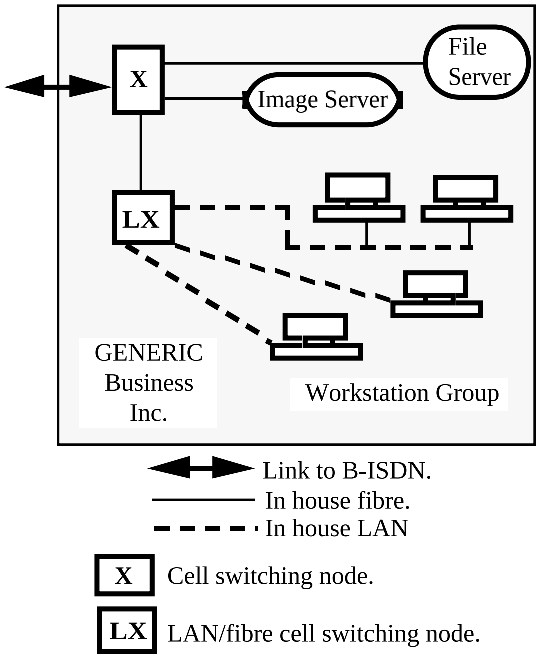

Figure 5.

The remainder of this paper will discuss the implementation of Figure 5 - that present LAN technologies, data link layers in their own rights, be put to use as the physical layer of an ATM cell transport system.

4.1 Ethernet encapsulated ATM.

One widespread LAN technology is Ethernet. An ATM fibre may be emulated between any two Ethernet interfaces that can directly exchange packets. For conceptual simplicity we encapsulated ATM cells within packets as shown in Figure 6 [20,21]. Using a unique EtherType code for ATM-over-Ethernet packets (we chose 0xEEEE) enables our ATM traffic to co-exist on the LAN with more traditional traffic (eg. IP).

Figure 7(a) shows the modified cell switch - consisting of a standard switch fabric with one fibre port, and one or more Ethernet ports. Outgoing cells are encapsulated and transmitted directly to the Ethernet interface of the destination endpoint AAL. Incoming cells are extracted from their packets and passed through the switch fabric without further processing.

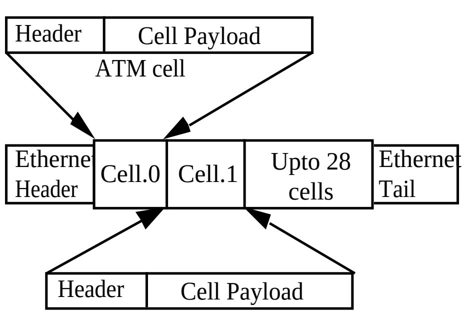

Figure 6

Figure 7(b) shows the emulated endpoint. An Ethernet interface is overlaid with an ATM layer that extracts cells from incoming packets before demultiplexing cells based on their VPI/VCI. Cells coming from the AALs above are placed into Ethernet packets and sent directly back to the cell switch.

4.2 Functions of the LAN/fibre switching node.

Creating a virtual connection involves mappings within the switch to associate (Input Port, VPI in, VCI in) with (Output Port, VPI out, VCI in). A LAN/fibre switching node behaves in a very similar way - The difference being that the output information is expanded to (Ethernet port, Destination Address, VPI out, VCI out). The inherent low speed of the LAN sections of the network mean this switch node may even be be based on a PC or workstation

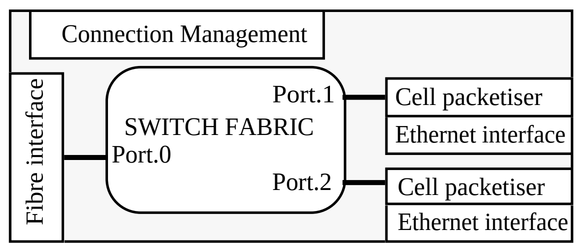

(a) LAN/fibre Switch Node

Figure 7(a)

with suitable interfaces - modifying cell headers, encapsulating cells, and extracting cells in software. Our test network used a SPARCStation1 as the switch node. The switch software was implemented with a kernel resident STREAMs module, with Ethernet access provided through the Network Interface Tap (NIT). An in-house

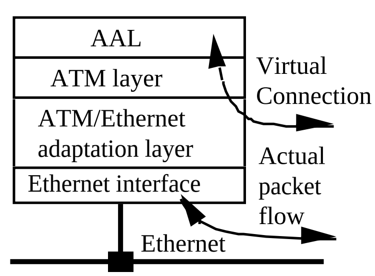

(b) End user workstation

Figure 7(b)

signalling protocol enabled ‘attached’ nodes to establish and remove virtual connections on-demand [22]. It also provided for boot-time discovery of the switch node’s Ethernet address by end-user nodes.

It is essential that at least one LAN/fibre switch node exists, in order to connect the in-house ‘ATM’ network with the external B-ISDN. This switch node will also be responsible for exchanging connection requests between the in-house network and the B-ISDN, performing signalling protocol conversions if necessary.

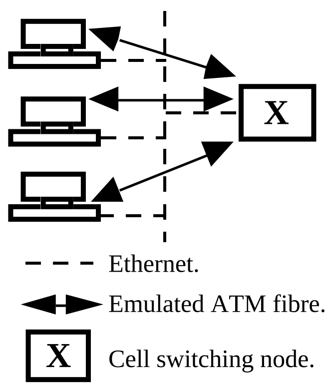

Figure 9

4.3 The local ATM node.

The emulated ATM endpoint in Figure 7 does not clearly show how different services will be supported in practice. Figure 8 shows a possible architecture that supports data and video/audio virtual connections. ATM services might be provided in software once we have accepted that ‘speed of operation’ is not a major criterion and suitably powerful machines are used.

The goal of Figure 8 is to allow easy creation of an “ATM interface”. The ATM layer and AAL services are provided as additional kernel or OS resident drivers, sharing the Ethernet interface with any other networking services in use. The major limitation is that video/audio traffic is processed by the workstation’s cpu. Unless the workstation is very powerful, this will severely limit the quality of multi-media service that is possible.

Alternative implementations may offload some or all of the video/audio processing to dedicated hardware codecs, with the main cpu simply providing the ATM layer (and possibly the AAL1 or AAL2 service) if necessary.

5. Implementation possibilities.

We have implemented a small ATM-over-Ethernet network. Four Sun SPARC based workstations were used to create a network of 3 ATM endpoints and 1 cell switch node. STREAMs modules were used to provide the AAL and ATM/Ethernet interface, with the SunOS Network Interface Tap (NIT) providing access to the Ethernet itself [21].

Our first network supported only AAL3/4 and AAL5 connections, for services such as TCP/IP over ATM. This required no extra multi-media hardware nor software codecs. All four machines shared the same Ethernet cable, emulating three separate fibres connecting each machine individually to the switch node (Figure 9). This was achieved by using directed Ethernet transmissions. The cell switch would encapsulate cells for a particular node in packets sent directly to the Ethernet address of that node. An in-house signalling system [22] was developed that provided ‘on-demand’ creation and removal of virtual connections.

As the ultimate goal is to incorporate the services of a hardware video/audio codec, we addressed the issue of processing video and audio traffic. Having decided that such traffic must not be processed by the workstation CPU, and should not traverse the system bus, two possibilities appeared. Figure 10 shows the first one, where no additional hardware is added to the workstation, and the signalling entity remains entirely within the workstation. The codec is set up with its own Ethernet interface and hardware (or software) capable of providing AAL1 or AAL2 services. An auxiliary link (e.g. serial line) conveys information from the signalling entity to the codec system to inform it of the VPIs and VCIs being used for any given video connection. This scheme requires the signalling system and switch node to be aware that the

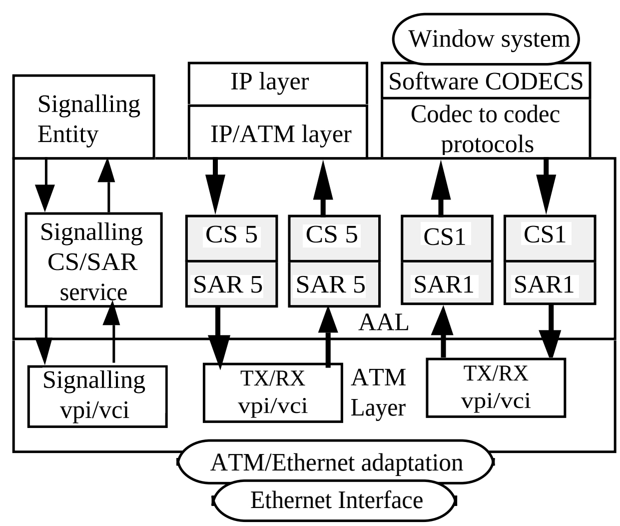

Figure 8

multi-media terminal has different Ethernet addresses for video and data ATM traffic.

However we wanted an 'integrated' terminal with all services converging on the same Ethernet address. This required the development of an external network interface (a 'splitter') that accepted all traffic on a single Ethernet port, and rerouted video or data traffic to either the codec hardware or the Sun workstation [21, 23, 27]. Figure 11 shows the 'splitter' taking the task of encapsulating ATM cells in packets, and extracting cells from packets. A dedicated link is used to carry cells from the 'splitter' to

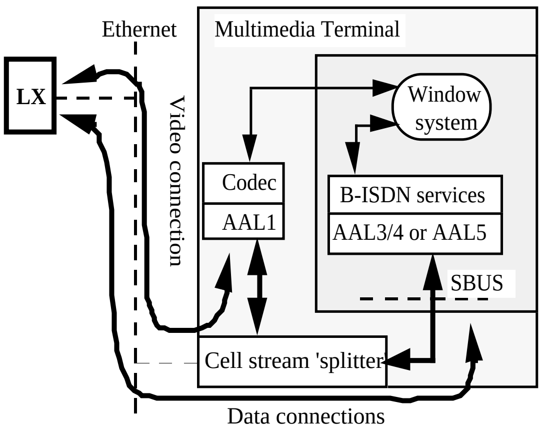

Figure 11

the codec. Future development of the 'splitter' will involve the addition of full ATM layer (and some AAL) services in hardware.

The network in Figure 9 was tested using two SPARCStation1+ workstations configured as shown in Figure 11. A custom designed interface based around Xilinx programmable logic arrays provided the 'splitter' function. The codecs were CLI Rembrandts running at 384 kbit/sec each way, using a proprietory coding scheme. An external processor board provided a form of AAL1 service to the codec. The video was displayed on the workstations screens using a RasterOps SPARC Card TC/PIP. The third station, a Sun IPC, contained only a software AAL3/4 and AAL5 interface to support TCP/IP over ATM. A SPARCStation1 was used to provide the ATM cell switch node. All 4 machines were running SunOS 4.1.1.

The system was tested with two virtual connections carrying video and audio, one in each direction, and TCP/IP connections between all 3 machines. All virtual connections were established through the switch 'on demand', rather than using preconfigured VPI/VCI's. Due to other demands being made on our codecs, there was no time to perform precise measurements of the networks performance limits. However the colour picture quality was quite acceptable, as was the end to end delay (subjectively around 1 second).

This arrangement did show that ATM services may be

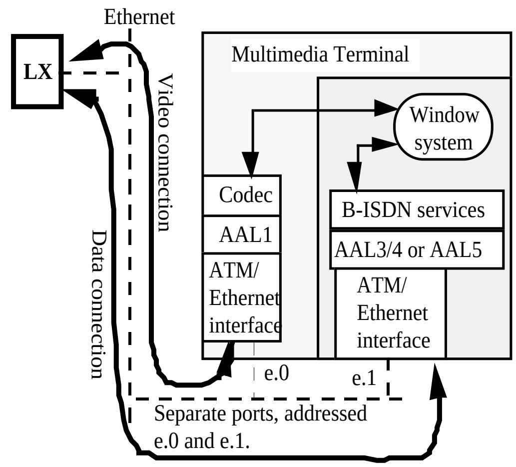

Figure 10

introduced into a computing environment and simultaneously support multimedia and non-multimedia terminals. The first phase of an office ATM evolution might simply be the addition of a LAN/ATM-fibre gateway and appropriate AAL3/4 or AAL5 software to selected workstations and PC's. Multimedia services may then be added as hardware or software codecs become available. As noted in section 4.1, ATM traffic may share the office LAN with other protocols during an overlap period.

6. Potential cell transport performance.

Encapsulation of ATM cells within LAN packets provides ATM connectivity rather than multi-Mbit/sec speeds. There is some tradeoff here, as packing more cells into a LAN packet reduces overheads but increases the packet processing time.

6.1 Theoretical packet use efficiency.

Assume each terminal has its own physical Ethernet connection to the switch. Figure 12 shows a graph of peak cell rate against the number of cells per packet, using the encapsulation scheme of Figure 6. Each cell is 424 bits long, the packet overhead (header, trailer, preamble, and inter packet delay) constitute 304 bit times, and a bit takes 0.1 µs. On this basis a full packet of 28 cells is 1217.6 µs long. One packet every 1217.6 µs corresponds to 22996 cells per second.

For a given average bitrate, the delay between cell arrival bursts increases in proportion to the number of cells in a packet. For example a minimum delay of 1217.6 µS will be experienced by cells being transported in groups of 28. This delay needs to be considered when running interactive connections such as video or voice.

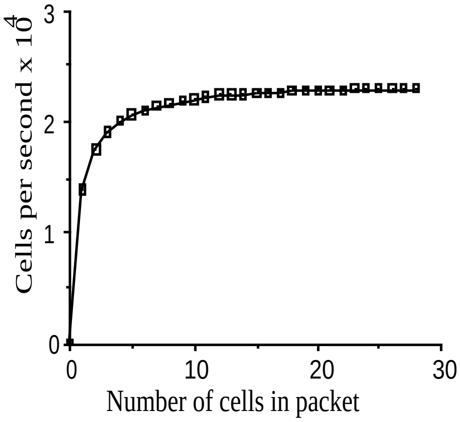

Figure 12

Figure 12 shows that the theoretical cell rate does not increase significantly when more than 10 cells are placed in each packet - suggesting that there is little to be gained by striving to maximise the number of cells per packet. In fact reduced cell packing reduces the buffering needed by codecs to smooth out bursty cell arrivals. It also improves the sharing of the medium, and reduces impact of packet loss on cell loss.

Despite the variables added by Ethernet cell transport, people should not be deterred from establishing video connections, as such services are already being used across the Internet today [7], a service that makes no claims to transferring data reliably.

6.2 Alternative cell transport schemes.

It is possible to encapsulate ATM cells differently, as shown in Figure 13. It assumes that the header of cells associated with a given AAL_SDU will remain fixed. The first cell of a burst is carried as a whole, and only the payloads of subsequent cells are carried. Large AAL_SDUs will save quite a bit of overhead, allowing the equivalent of 31 cells to be transported in each Ethernet packet.

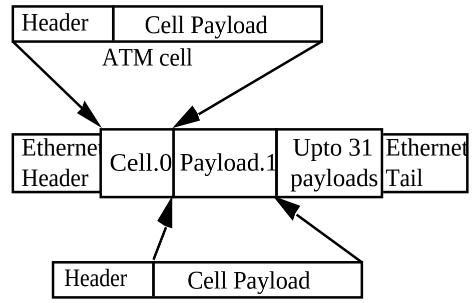

Figure 13

At the receiving end the cell stream is reconstructed by prepending the first header onto all the payloads before passing up to the ATM layer. However, this scheme assumes there are sequences of cells with common cell headers. This is not true if the ATM layer is multiplexing cells from multiple AAL_SDUs. In addition, AAL5 modifies the header of the last cell from an AAL_SDU - necessitating two packets per AAL_SDU even under the most ideal conditions. Finally this mechanism would fragment the cell/packet stream whenever a switch node modified the CLP bit in the header. The real nature of ATM traffic quickly removes any packet use efficiency gains this scheme might offer.

More complex schemes are possible, for example the entire AAL could be distributed - with the workstations only containing the Convergence Sublayers and the LAN/fibre switching nodes executing the Segmentation and Reassembly functions on traffic going from the internal network to the wider B-ISDN. This would, however, impose significant complexity and processing load on the LAN/fibre switch - something we wish to avoid.

6.3 Performance examples from a software AAL.

TCP/IP traffic could be carried across AAL5 virtual connections at rates between 500kbit/sec and 600 kbit/sec, using the NIT to access the Ethernet. Raw virtual connections through a cell switch were capable of one way traffic up to 1 Mbit/sec using 1200 byte AAL Service Data Units (AAL_SDUs). Using 4000 byte AAL_SDUs resulted in throughput around 1.5Mbit/sec (Our AAL3/4 emulation was always slower, as the processing and SAR layer overheads were higher). The efficiency of transmission decreased as AAL_SDUs became smaller. Connections carrying 200 byte AAL_SDUs would achieve upwards of 400 kbit/sec.

A significant amount of time was spent performing ATM layer and AAL processing in the STREAMs stack. (STREAMs has been shown to have very poor performance by other researchers, e.g. McCanne and

Jacobson during their development of the “BSD Packet Filter” [25]).

Speed tests using an internal software loopback connection under the AAL only achieved between 13% and 20% improvement over connections through the Ethernet/ATM link. This showed that the figures we obtained were not due to limitations of the LAN media. They do show, however, how important the development of hardware AALs will be.

7. Conclusion.

Broadband ISDN offers a unified network interface to vast quantities of information. ATM provides the flexibility in data transport, and the fibre provides raw transport speeds needed to bring images and sounds to the desktop. However we have seen that new image compression techniques have reduced bandwidth requirements to the sub-Mbit/sec range.

We have proposed the building small ATM networks using traditional wire based LANs to transport cells, rather than fibre. The speed penalty is considered a price worth paying in order to enable desktop workstations to create direct ATM connections to remote databases and information sources in the wider B-ISDN. A network architecture is proposed consisting of at least one LAN/fibre switch node whose job is to switch cells between inhouse physical LAN links and an external B-ISDN fibre coming to the premises. ATM terminals are created by adding software AALs and ATM layers utilising LAN interfaces that typically already come with many workstations. The evolutionary path to full ATM will be able to occur without disruption of ATM applications that have been installed on the workstations in the interim period.

Ackowledgement.

The assistance of AOTC through Telecom Research Laboratories research contract 7064 is hereby acknowledged.

References.

[1] Shoshana Loeb, “Delivering Interactive Multimedia Documents over Networks”, IEEE Communications Magazine, May 1992, p. 53.

[2] K. Wallace, “The JPEG Still Picture Compression Standard”, Communications of the ACM, Vol. 34, No.4, April 1991.

[3] Jonathan Rosenberg, Robert E. Kraut, Louis Gomez, C. Alan Buzzard, “Multimedia Communications for Users”, IEEE Communications Magazine, May 1992, p.26.

[4] D. Le Gall, “MPEG: A Video Compression Standard for Multimedia Applications”, Communications of the ACM, Vol. 34, No. 4, April 1991.

[5] Joe Sutherland, Larry Litteral, “Residential Video Services”, IEEE Communications Magazine, July 1992, p. 37.

[6] M. Liou, “Overview of the px64 kbit/s Video Coding Standard”, Communications of the ACM, Vol. 34, No. 4, April 1991.

[7] Thierry Turletti, “H.261 software codec for videoconferencing over the internet”, INRIA Technical Report RR1834, January 1993.

[8] CCITT SGXVIII - Report R109, Draft recommendation I.150, "B-ISDN Asynchronous Transfer Mode Functional Characteristics", July 1992.

[9] Martin de Prycker, “Asynchronous Transfer Mode - Solution for Broadband ISDN”, Ellis Horwood, 1991.

[10] CCITT SGXVIII Draft recommendation I.362, "B-ISDN ATM Adaptation Layer (AAL) functional description", June 1992.

[11] CCITT SGXVIII Draft recommendation I.363, "B-ISDN ATM Adaptation Layer (AAL) specification", June 1992.

[12] CCITT SGXVIII Draft recommendation I.363, "B-ISDN ATM Adaptation Layer (AAL) specification", January 1993, Section 6.

[13] ANSI T1S1.5/91-449, “AAL5 - A New High Speed Data Transfer AAL.”, ANSI T1S1, November 1991.

[14] Apple Computer, Bellcore, Sun Microsystems, Xerox, “Network Compatible ATM for Local Network Applications, Phase 1”, Version 1.0, April 1992.

[15] Jon Postel, “Internet Protocol - DARPA Internet Program Protocol Specification”, RFC-791, September 1981, Section 1.4.

[16] D. Comer, “Internetworking with TCP/IP - Principals, Protocols, and Architecture”, Prentice Hall, 1988.

[17] CCITT SGXVIII Revised recommendation I.361, "B-ISDN ATM Layer specification", June 1992.

[18] CCITT SGXVIII - Report R111, Draft recommendation I.311, “B-ISDN General Network Aspects”, July 1992.

[19] “The Ethernet - A Local Area Network”, Version 1.0, Digital Equipment Corporation, Intel Corporation, Xerox Corporation, September 1980.

[20] G.J. Armitage, K.M. Adams, “Architecture of a Multimedia Desktop Workstation”, Proc. Australian Video Communications Workshop, Melbourne, July 1990, p.76.

[21] G.J. Armitage, K.M. Adams, “Prototyping an ATM Adaptation Layer in a Multi-Media Terminal”, Int. J. of Digital and Analog Communication Systems, Vol.6, No.1, January 1993, pp.3-14.

[22] G.J. Armitage, K.M. Adams, “Towards Shared ATM Fibres, A Pseudo Metasignalling Protocol, and ATM Lans.”, Proc. Australian Broadband Switching and Services Symposium, Melbourne, July 1992, p.271.

[23] K.K.L. Wong, K.M. Adams, “Hardware Realization of an ATM Cell Splitter in a B-ISDN Terminal”, Proc. Australian Video Communications Workshop, Melbourne, July 1990, p.86.

[24] Shoshana Loeb, “Delivering Interactive Multimedia Documents over Networks”, IEEE Communications Magazine, May 1992, p. 56.

[25] Steve McCanne, Van Jacobson, “The BSD Packet Filter: A New Architecture for User-level Packet Capture”, December 1992 preprint of submission to 1993 Winter USENIX conference, January 1993, San Diego, CA.

[26] Juha Heinanen, “Multiprotocol Encapsulation over ATM Adaptation Layer 5”, IETF Network Working Group, DRAFT Request For Comment, February 1993

[27] S. Liew, K.M. Adams, “Hardware Design Aspects in Implementing ATM Functions in an Multi-media Workstation”, Proc. 2nd Australian Multi-media Communications, Applications, and Technology Workshop, Melbourne, 13-14th July 1992, p.189.

Proc IEEE Computer Society 18th Conference on Local Computer Networks, Minneapolis, MN, September 19-22, 1993. This pdf reconstructed Dec 2005.